

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

PATENT APPLICATION

in accordance with Art. 158(3) EPC

(43) Date of publication:
03.01.1996 Bulletin 1996/01

(21) Application number: 95902965.3

(22) Date of filing: 15.12.1994

(51) Int. Cl.⁵: G06F 9/38

(86) International application number: PCT/JP94/02112

(87) International publication number: WO 95/16954
(22.06.1995 Gazette 1995/26)

(84) Designated Contracting States:
DE FR GB

(30) Priority: 15.12.1993 US 168744

(71) Applicants:
• SILICON GRAPHICS, INC.
Mountain View, CA 94039-7311 (US)
• KABUSHIKI KAISHA TOSHIBA
Kawasaki-shi, Kanagawa-ken 210, Tokyo (JP)


(72) Inventors:
• JOSHI, Chandra, S.
Saratoga, CA 95070 (US)
• RODMAN, Paul
Palo Alto, CA 94306 (US)
• HSU, Peter, Yan-Tek
Fremont, CA 94555 (US)
• NOFAL, Monica, R.
Los Altos, CA 94024 (US)

(74) Representative: Ritter und Edler von Fischern,
Bernhard, Dipl.-Ing. et al
D-81925 München (DE)

(54) APPARATUS FOR PROCESSING INSTRUCTION IN COMPUTER SYSTEM

(57) A computer system includes first and second instruction storage circuits, each of which stores N instructions for parallel output. An instruction dispatch circuit connected to the first instruction storage circuit dispatches L instructions stored in the first instruction storage circuit, and L is hereby equal to, or smaller than N. An instruction load circuit connected to the first and second instruction storage circuits loads the L instructions from the second instruction storage circuit to the first instruction storage circuit after the L instructions are dispatched from the first instruction storage circuit and before another instruction is dispatched from the first instruction storage circuit. An instruction memory stores a plurality of strings of instructions and a branch memory stores a plurality of branch anticipation entries.

FIG.2



Description

Technical Field

5 This invention relates to computing systems and, more particularly, to an apparatus for processing instructions in a computing system.

Background Art

10 In a typical computing system, instructions are fetched from an instruction memory, stored in a buffer, and then dispatched for execution by one or more central processing units (CPU's). Figs. 10A-10C show a conventional system where up to four instructions may be executed at a time. Assume the instructions are alphabetically listed in program sequence. As shown in Fig. 10A, an instruction buffer 10 contains a plurality of lines 14A-C of instructions, wherein each line contains four instructions. The instructions stored in buffer 10 are loaded into a dispatch register 18, comprising four
 15 registers 22A-D, before they are dispatched for execution. When four instructions are dispatched simultaneously from dispatch register 18, then four new instructions may be loaded from buffer 10 into dispatch register 18, and the process continues. However, sometimes four instructions cannot be dispatched simultaneously because of resource contention or other difficulties. Fig. 10B shows the situation where only two instructions (A,B) may be dispatched simultaneously. In known computing systems, the system must wait until dispatch register 18 is completely empty before any further
 20 instructions may be transferred from buffer 10 into dispatch register 18 to accommodate restrictions on code alignment and type of instructions that may be loaded at any given time. Consequently, for the present example, at most only two instructions (C,D) may be dispatched during the next cycle (Fig. 10C), and then dispatch register 18 may be reloaded (with instructions E,F,G, and H). The restriction on the loading of new instructions into dispatch register 18 can significantly degrade the bandwidth of the system, especially when some of the new instructions (e.g., E and F) could have
 25 been dispatched at the same time as the instructions remaining in the dispatch register (C,D) had they been loaded immediately after the previous set of instructions (A,B) were dispatched.

Another limitation of known computing systems may be found in the manner of handling branch instructions where processing continues at an instruction other than the instruction which sequentially follows the branch instruction in the instruction memory. In the typical case, instructions are fetched and executed sequentially using a multistage pipeline.
 30 Thus, a branch instruction is usually followed in the pipeline by the instructions which sequentially follow it in the instruction memory. When the branch condition is resolved, typically at some late stage in the overall pipeline, instruction execution must be stopped, the instructions which follow the branch instruction must be flushed from the pipeline, and the correct instruction must be fetched from the instruction memory and processed from the beginning of the pipeline. Thus, much time is wasted from the time the branch condition is resolved until the proper instruction is executed.

35 It is an object of the present invention to provide an apparatus for processing instructions in a computing system wherein four instructions are always made available for dispatching regardless of how many instructions are previously dispatched, and without regard to code alignment or instruction type.

In one embodiment of the invention, a computing system has first and second instruction storing circuits, each instruction storing circuit storing N instructions for parallel output. An instruction dispatch circuit, coupled to the first
 40 instruction storing circuit, dispatches L instructions stored in the first instruction storing circuit, wherein L is less than or equal to N. An instruction loading circuit, coupled to the instruction dispatch circuit and to the first and second instruction storing circuits, loads L instructions from the second instruction storing circuit into the first instruction storing circuit after the L instructions are dispatched from the first instruction storing circuit and before further instructions are dispatched from the first instruction storing circuit.

45 It is another object of the present invention to provide an apparatus for processing instructions in a computing system wherein branches are predicted at the time of instruction fetch, and the predicted target instruction is fetched immediately so that the target instruction is available for execution immediately after the branch instruction is executed.

Disclosure of Invention

50 In one embodiment of this aspect of the invention, an instruction memory stores a plurality of lines of a plurality of instructions, and a branch memory stores a plurality of branch prediction entries, each branch prediction entry containing information for predicting whether a branch designated by a branch instruction stored in the instruction memory will be taken when the branch instruction is executed. Each branch prediction entry includes a branch target field for indicating
 55 a target address of a line containing a target instruction to be executed if the branch is taken, a destination field indicating where the target instruction is located within the line indicated by the branch target address, and a source field indicating where the branch instruction is located within the line corresponding to the target address. A counter stores an address value used for addressing the instruction memory, and an incrementing circuit increments the address value in the counter for sequentially addressing the lines in the instruction memory during normal sequential operation. A counter

by the branch instruction stored in the instruction memory will be taken when the branch instruction is executed. That way the line containing the target instruction may be fetched and entered into the pipeline immediately after the line containing the branch instruction. An invalidate circuit invalidates any instructions following the branch instruction in the line containing the branch instruction and prior to the target instruction in the line containing the target instruction.

In order to accomplish the above objects, an apparatus according to the present invention for processing instruction in a computing system comprises first and second instruction storing circuits, each instruction storing circuit storing N instructions for parallel output, an instruction dispatch circuit, coupled to the first instruction storing circuit, for dispatching L instructions stored in the first instruction storing circuit, wherein L is less than N, and an instruction loading circuit, coupled to the instruction dispatch circuit and to the first and second instruction storing circuits, for loading L instructions from the second instruction storing circuit into the first instruction storing circuit after the L instructions are dispatched from the first instruction storing circuit and before further instructions are dispatched from the first instruction storing circuit.

Also, another apparatus according to the present invention for processing instruction in a computing system comprises an instruction storing circuit for storing N instructions for parallel output, an instruction dispatch circuit, coupled to the instruction storing circuit, for dispatching L instructions stored in the first instruction storing circuit, wherein L is less than N, and an instruction queue for storing M lines of N instructions from an instruction memory, an instruction loading circuit, coupled to the instruction storing circuit and to the instruction queue, for loading L instructions from the instruction queue into the instruction storing circuit after the L instructions are dispatched from the instruction storing circuit and before further instructions are dispatched from the instruction storing circuit.

Also, a further apparatus according to the present invention for processing instruction in a computing system comprises an instruction memory for storing a plurality of lines of a plurality of instructions and a branch memory for storing a plurality of branch prediction entries, each branch prediction entry containing information for predicting whether a branch designated by a branch instruction stored in the instruction memory will be taken when the branch instruction is executed.

In accordance with the present invention, four instructions are always made available for dispatching regardless of how many instructions are previously dispatched, and without regard to code alignment or instruction type.

Also, branches are predicted at the time of instruction fetch, and the predicted target instruction is fetched immediately so that the target instruction is available for execution immediately after the branch instruction is executed.

Brief Description of Drawings

Fig. 1 is a block diagram showing instruction fetch and dispatch in a particular embodiment of a computing system according to the present invention;

Fig. 2 is a block diagram of a particular embodiment of an apparatus according to the present invention for fetching and dispatching instructions;

Fig. 3 is a block diagram illustrating the operation of the instruction queuer of Fig. 3;

Fig. 4 is a block diagram of an alternative embodiment of an apparatus according to the present invention for fetching and dispatching instructions;

Fig. 5 is a block diagram of a particular embodiment of an apparatus according to the present invention for predicting branches;

Fig. 6 is a block diagram of a particular embodiment of an entry in the branch cache shown in Fig. 5; and

Fig. 7 is a block diagram showing the Fetch (F) stage of the instruction pipeline according to the present invention.

Fig. 8 is a block diagram showing the Decode (D) and Address Generation (A) stages of the instruction pipeline according to the present invention.

Fig. 9 is a block diagram showing the Execute (E) and Writeback (W) stages of the pipeline according to the present invention.

Fig. 10 is a block diagrams showing instruction fetch and dispatch in a known computing system;

Best Mode for Carrying Out the Invention

Figs. 1A-D are block diagrams showing instruction fetch and dispatch in a particular embodiment of a computing system according to the present invention. As in the example shown in Figs. 10A-D, assume two instructions (A,B) are dispatched initially. However, unlike the example in Figs. 10A-D, the two dispatched instructions (A,B) are immediately replaced by the next two sequential instructions (E,F) as shown in Fig. 1B. Thus, four instructions are available for dispatch in the next clock cycle. A pointer 26 is used to keep track of which instruction follows the previously dispatched instructions in the program sequence. If three instructions are dispatched in next clock cycle, as shown in Fig. 1C, then the instruction indicated by pointer 26, together with the two sequentially following instructions, may be released by

enabling the appropriate registers 22A, 22C, and 22D. Immediately thereafter, dispatch register 18 is loaded with the next three instructions in the program sequence from instruction buffer 10.

It should be noted at this point that two lines in the instruction buffer may need to supply the instructions loaded into dispatch register 18. For example, line 14C supplies instructions (G,H) and line 14B supplies instruction (I) to dispatch register 18 in Fig. 1C. Assuming four instructions per line, the line which contains the next sequential program instruction to be loaded into dispatch register 18 may be termed the "leading quad", and any next further line which simultaneously supplies instructions to dispatch register 18 to complete the fill operation may be termed the "trailing quad". When the leading quad is emptied by the fill operation, then the contents of the buffer may be advanced by one line as shown in Fig. 1D. In Fig. 1D, two more instructions (F,G) are dispatched, and two instructions (J,K) are loaded in their place.

Fig. 2 is a block diagram of a particular embodiment of an apparatus 30 for fetching and dispatching instructions according to the present invention. Apparatus 30 includes an instruction cache 34 which stores a plurality of lines of instructions that may be addressed by an address value received on a communication path 38. In this embodiment, each line stores four 32-bit instructions and communicates all the instructions in a line to a predecode circuit 42 over a communication path 46. Predecode circuit partially decodes the four instructions and communicates the four partially decoded instructions to an instruction queuer 50 over a communication path 54 and to dispatch multiplexers 58A-D over a queue bypass path 62.

Instruction queuer 50 includes four queue sections 66A-D, one for each instruction in each line. All four queue sections have the same construction, so only the details of queue section 66A shall be described. Queue section 66A includes a plurality, e.g., six, serially connected instruction buffers IBUF0-IBUF5. Each instruction buffer is coupled to a multiplexer 68 through a corresponding multiplexer input path 70A-F. Multiplexer 68 selects one of the instructions from among instruction buffers IBUF0-IBUF5 in response to signals received over a line 72A and communicates the selected instruction to a dispatch multiplexer 58A over a communication path 74A. The current instruction in register 22A is also communicated to the input of dispatch multiplexer 58A over a feedback communication path 76D. Dispatch multiplexer 58A thus selects from among the output of multiplexer 68, queue bypass path 62, or feedback communication path 76A in response to signals received over a Q0MXSEL line 80A to communicate an instruction to register 22A over a communication path 82A. Register 22A then loads the received value in response to clock signals applied to the register (clocked registers are indicated by the angled symbol on the left side of each register) and then dispatches the instruction when possible.

Queue sections 66B-D also select instructions within one of their serially connected buffer sections in response to signals received over lines 72B-D, respectively, and communicate the selected instructions to dispatch multiplexers 58B-D over respective communication paths 74B-D. Dispatch multiplexers 58B-D communicate instructions, selected by signals received over Q1MXSEL-Q3MXSEL lines, to their respective registers 22B-D over communication paths 82B-D.

Apparatus 30 selects which instructions are to be presented to dispatch register 18 in the following manner. The first time a line of instructions is retrieved from instruction cache 34, instruction queuer 50 is empty, and multiplexers 58A-D select the instructions from queue bypass path 62. Instructions are then dispatched, and a new line of instructions are read from instruction cache 34.

In general, a new line of instructions is read from instruction cache 34 on every clock cycle. If four instructions were dispatched every clock cycle, then dispatch register would always be loaded from queue bypass path 62. However, at any given cycle anywhere from zero to four instructions may be dispatched. Thus, if not all instructions are dispatched, then only certain ones of registers 22A-D are loaded from queue bypass path 62 pursuant to the number of instructions dispatched. The previously read line of instructions is then loaded into IBUF0 in each queue section 66A-D, and a new line of instructions is read from instruction cache 34. Thereafter, instructions are loaded from IBUF0 in the appropriate queue section 66A-D and from queue bypass path 62. For example, if two instructions are dispatched on the first cycle, then registers 22A-B are loaded from queue bypass path 62, registers 22C-D are reloaded with the same instructions via communication paths 76C-D, the previously read line of instructions is loaded into IBUF0 in queue sections 66A-D, and a new line of instructions is read from instruction cache 34. If only one instruction is dispatched during the next clock cycle, then register 22C is loaded from IBUF0 in queue section 66C, registers 22A, 22B, and 22D are reloaded with the same instructions via communication paths 76A, 76C, and 76D, the line of instructions stored in IBUF0 in each queue section 66A-D is advanced to IBUF1 in each queue section, the previously read line of instructions is loaded into IBUF0 in queue sections 66A-D, and a new line is read from instruction cache 34. The lines of instructions are advanced within queue sections 66A-D until the buffer is full. At that time the apparatus stalls further loading of instruction lines into the queue. This manner of operation allows the instruction prefetch operation to be decoupled from the dispatch operation.

A RDPTR register 86 stores a value I STATE [4:0] for controlling the operation of instruction queuer 50. STATE [4:2] is used to determine which buffer IBUF0-IBUF5 in each queue section 66A-D supplies the next instruction to registers 22A-D, and STATE [1:0] functions as pointer 26 in Figs. 1A-1C (a modulo-4 counter) to indicate which instruction is to be dispatched next. An F INST register 90 stores an INST CONSUME value indicating how many instructions are consumed in every cycle (i.e., the sum of queuer register clock enables, or the total number of instructions dispatched from dispatch register 18 whether valid or not). The INST-CONSUME value is discussed in conjunction with Fig. 8. The INST CONSUME value is added to STATE [4:0] by an adder 92 to point to the next instruction to be dispatched. STATE [4:2]

The updated value of STATE [4:0] is loaded back into RDPTR register 86 and communicated to a queuer mux select circuit 98 over communication paths 99A and 99B. If STATE[4:2]="101" (=5); the instruction buffer is full, and the apparatus stalls further loading of instruction lines into the queue.

Queuer mux select circuit 98 presents the next four sequential instructions (in program order) to dispatch register 18 in accordance with the values of STATE [4:2] and STATE [1:0]. Fig. 3 and Table 1 show which buffer in each queue section 66A-D supplies the next instruction to its corresponding register 22A-D in dispatch register 18 for the different values of STATE [1:0].

TABLE 1

STATE[1:0]	Q0MXSEL	Q1MXSEL	Q2MXSEL	Q3MXSEL
0	STATE[4:2]	STATE[4:2]	STATE[4:2]	STATE[4:2]
1	STATE[4:2]-1	STATE[4:2]	STATE[4:2]	STATE[4:2]
2	STATE[4:2]-1	STATE[4:2]-1	STATE[4:2]	STATE[4:2]
3	STATE[4:2]-1	STATE[4:2]-1	STATE[4:2]-1	STATE[4:2]

Thus, if STATE[1:0]=2 and STATE[4:2]=3, then registers 22C and 22D will be presented with the last two instructions in the leading quad (IBUF3), and registers 22A and 22B will be presented with the first two instructions in the trailing quad (IBUF2).

The described apparatus for fetching and dispatching instructions may be used in many environments with or without modification. For example, assume integer, memory, and floating point instructions are stored in instruction cache 34, and they may be mixed within a line of instructions. If there is a problem with resource contention and data dependencies with an instruction or type of instruction (e.g., floating point instructions), then those instructions may be dispatched into another queue where they can wait for the resource contention and data dependencies to clear without holding up dispatching of the other instructions.

Fig. 4 is a block diagram of an alternative embodiment of an apparatus 104 according to the present invention for fetching and dispatching floating point instructions that may have been previously dispatched from dispatch register 18 in Fig. 4. From inspection it is apparent that apparatus 104 operates much like apparatus 30 in Fig. 2, except apparatus 104 also provides for storing data together with the instructions to handle integer store operation data or floating point register data that is to be loaded from the integer register.

The previously described apparatus also facilitates processing instructions in a computing system according to the present invention wherein branches are predicted at the time of instruction fetch, and wherein the predicted target instruction is fetched immediately so that the target instruction is available for execution immediately after the branch instruction is executed. Fig. 5 is a block diagram of a particular embodiment of an apparatus 110 according to the present invention for predicting branches. A branch prediction cache 114 is used to predict the outcome of branch instructions stored in instruction cache 34. For example, instruction cache 34 may be a 16KB direct-mapped cache which outputs four instructions per cycle as noted above. In this embodiment, branch prediction cache 114 is also direct mapped and may contain 1K entries (one entry per four instructions in instruction cache 34). Instruction cache 34 and branch cache 114 are accessed in parallel in the fetch stage of the pipeline through communication path 38 which receives an index (address) value from a counter 116. Of course, instruction cache 34 and branch prediction cache 114 could be accessed with different addresses if desired.

Fig. 6 shows a sample entry 120 from branch prediction cache 114 and an example of branch prediction. Entry 120 includes a valid field 124 for predicting whether the branch is taken (0=not predicted; 1=predicted), an index field 128 which is the instruction cache index of the branch target instruction, a source field (SRC) 132 which indicates the position of the last instruction to be executed within the line containing the branch instruction, and a destination field (DST) 134 which indicates the position of the branch target instruction within the line fetched by the cache index.

In this embodiment, each branch instruction actually comprises two instructions. The first instruction, termed the initial branch instruction, computes the branch target and the branch condition. The second instruction, termed a delay instruction, immediately follows the initial branch instruction and is used to actually change the program flow to the branch target instruction. Consequently, the source field 132 typically indicates the position of the delay instruction within the instruction line as shown in Fig. 6.

The address value in counter 116 is communicated to an incrementing circuit 138 which increments the counter value by four (since there are four instructions per line) and communicates the incremented value to a multiplexer 142 through a communication path 144. Additionally, the value in the index field 128 of the branch cache entry is communicated to multiplexer 142 over a communication path 148. The value in the valid field 124 may be used to control the operation of multiplexer 142. Thus, if the branch is predicted ($V=1$), then instruction cache 34 will be addressed with the value from index field 128 in the next cycle. If the branch is not predicted ($V=0$), then instruction cache 34 will be addressed with the next sequential line of instructions as determined from incrementing circuit 138.

The value in source field 132 is communicated to a valid mask 148 through an OR circuit 150. If the branch is predicted, valid mask 148 invalidates all instructions in the current line which occur after the delay instruction associated with the branch, since they would not be executed if the branch were taken. For example, if the delay instruction is the third instruction in the line as shown in Fig. 6, then the fourth instruction will be invalidated. During the next clock cycle, the line (including any invalidated instructions) are communicated to instruction queuer 50 and queue bypass path 62 (Fig. 2), the value of the destination field is loaded into a register 152, the value of counter 116 is loaded with the value from index field 128, and instruction cache 34 is addressed to fetch the line which contains the predicted branch target instruction. The destination field in register 152 is then communicated to valid mask 148 through OR circuit 150 to invalidate the instructions which occur before the branch target instruction in the line. For example, if the branch target instruction is the second instruction in the line, then valid mask 148 invalidates the first instruction in the line. The line is then communicated to instruction queuer 50 and queue bypass path 62.

In this embodiment, all branch prediction cache entries are initialized with a valid field of zero (branch not predicted). When the program executes the first time, the result of each branch instruction is used to update the branch prediction cache entry (if necessary) by setting the valid bit to one, and by inserting the appropriate index, source, and destination values. Branch prediction thus may occur thereafter. If a branch previously taken is not taken at a later time, or if a branch not previously taken is taken at a later time then the branch cache entry is updated (and correct instruction fetched) accordingly (discussed below).

Additionally, dispatch register 18 breaks (holds) the superscalar instructions which occur after the delay instruction of a predicted branch in dispatch register 18 to avoid mixing target instructions with a current branch instruction. Furthermore, dispatch register 18 breaks (holds) the superscalar instructions at the second branch when two branches are stored in dispatch register 18 so that only one branch at a time is allowed to execute.

Figs. 7-9 are block diagrams of a particular embodiment of portions of an instruction pipeline according to the present invention showing how branch prediction operates. Where possible, reference numbers have been retained from previous figures. Instruction cache 34 may comprise an instruction memory and a tag memory as is well known in the art. The instruction memory portion may contain the lines of instructions, and the tag memory may contain the virtual address tags (and control information) associated with each line in the instruction memory. For the present discussion, only the tag memory portion (34A) of instruction cache 34 is illustrated. Tag memory 34A includes an application specific identification field (asid[7:0]), the instruction cache tag (tag[33:0], the high order 34 bits of the associated virtual address), a valid bit (V) and a region field (r[1:0]) for indicating the address space of the instruction.

Fig. 7 shows the Fetch (F) stage of the instruction pipeline. Counters 116A and 116B are the primary F stage program counter which addresses tag memory 34A and branch cache 114. The value in counter 116A (fpc[13:4]), which indexes a line in tag memory 34A, is communicated to tag memory 34A and to incrementing circuit 138 over communication path 38A. Incrementing circuit 138 adds one to the counter value and communicates the incremented value to multiplexer 142A and multiplexer 142B over communication path 144. Multiplexers 142A and 142B also receive the index field from branch cache 114 over communication path 148, and a correction address (described below) over a communication path 160. The value on communication path 160 (pc jam-bus[13:2]) is used to correct branch misprediction, cache misses, etc. Multiplexer 142B also receives a branch cache write address (bcwadr[13:4]) for updating the branch cache. The data used to update branch prediction cache 114 (bcwdata[14:0]) is communicated to a register 164 over a communication path 168. Multiplexers 142A and 142B select the appropriate address and communicate it to counters 116A and 116B, respectively.

A register 172 stores a parallel load bit (fpld) indicating whether counters 116A-B were loaded with the incremented value from incrementing circuit 138 or whether counters 116A-B were loaded from either communication path 148 or communication path 160, and a register 176 stores a value (fpc[3:2]) corresponding to the destination field of a branch prediction cache 114 entry (bits (4:3) of the bc (14:3) data on communication path 148). The values in registers 116A, 172, and 176 are combined with the output of tag memory 34A and stored in a queue register TBUF0, which is one of six registers (TBUF0-TBUF5) used to store tag data to correspond to the six instruction buffers IBUF0-IBUF5 in instruction queuer 50. Each register TBUF0-TBUF5 is coupled to multiplexers 180 and 184 which select the registers which correspond to the leading quad and trailing quad, respectively, in instruction queuer 50. The leading quad tag memory information is communicated to the next stage in the pipeline over a communication path 188, and the trailing quad tag memory information is communicated to the next stage in the pipeline over a communication path 190.

Fig. 8 shows the Decode (D) and Address Generation (A) stages of the instruction pipeline. In the D stage, bits (56:2) of the leading quad information from tag memory 34A is stored in a DLTAG register 200, and the trailing quad

information from tag memory 114 is stored in a DTTAG register 204. The destination field of the branch prediction cache 114 entry (fpc[3:2]), if any, associated with the leading quad information is communicated to a multiplexer 208. The other input to multiplexer 208 is coupled to an output communication path 210 of an adder 214 which contains the pointer value of the position of the next sequential instruction to be dispatched during normal sequential execution. Multiplexer 208 selects either the destination value or the next sequential instruction value and communicates the selected value to the output communication path 216 of DLTAG register 200. Communication path 216 is coupled to an input of a multiplexer 218.

The trailing quad tag memory information stored in DTTAG register 204 is communicated to multiplexer 218 and to a compare circuit 220 over a communication path 224. Multiplexer 218 selects the tag information corresponding to the first instruction to be executed next and outputs the selected information on a communication path 226 to an ATAG register in the A stage of the pipeline. The dispatch register pointer value is communicated to adder 214 over a communication path 228, the tag memory information is communicated to compare circuit 220 over a communication path 230, and the instruction cache index is communicated to a compare circuit 234.

Compare circuit 220 compares the leading quad tag to the trailing quad tag. If they do not match, then the leading quad instructions and the trailing quad instructions come from a different context, so they should not be dispatched simultaneously. A signal is provided on a communication path 238 to break the superscalar instructions when this occurs.

Compare circuit 234 compares the instruction cache index to the hex value "FFF" to determine if the end of the instruction cache is being addressed. If so, then it is desirable to break the superscalar instructions at the end of the cache line, and a signal is provided on a communication path 242 for that purpose.

Adder 214 receives a value indicating the sum of valid instructions dispatched over a communication path 250, and that value is used to increment the current dispatch register pointer value to produce the updated dispatch register pointer value on communication path 214.

During the D stage, register 90 (see also Fig. 2) is loaded with the value indicating the number of instructions consumed (both valid and invalid instructions), and this value is used to control the operation of instruction queuer 50 as discussed in conjunction with Fig. 2.

During the A stage, the actual branch address is generated. Since each branch instruction comprises an initial branch instruction followed by a delay instruction, and since the actual branch is accomplished after the delay instruction, the branch target address must be calculated relative to the delay instruction. Accordingly, when the tag information corresponding to the line containing the branch instruction is stored in ATAG register 227, a value indicating the relative position of the delay instruction within the line is selected by a multiplexer 250 and stored in a RELDLY register 254 via a communication path 258. The relative delay value is communicated to a branch target adder 260 over a communication path 264. Branch target adder also receives the ATAG register 227 value (which is the address of the first instruction in the line) via a communication path 268, and an offset value from an AOFFSET register 272 via a communication path 276. AOFFSET register 272 receives the 26-bit offset value from the branch instruction over a communication path 280, and subjects bits [17:2] of the offset value to a sign extension function in a sign extension circuit 284 (if necessary) prior to forwarding the offset value to branch target adder 260. AOFFSET register 272 also communicates the 26-bit offset value to a multiplexer 288 which also receives bits [27:2] of the branch target address calculated by branch target adder 260 over a communication path 292. Multiplexer 288 thus allows bits [27:2] of the calculated branch target address to be replaced by the offset value stored in AOFFSET register 272.

The output from branch target adder 260 is communicated to one input of a multiplexer 288. The other input to multiplexer 288 is a branch target address from a JUMP or JUMP REGISTER instruction received over a communication path 296 coupled to the general purpose register file. Thus, the selected branch target address will be the output from branch target adder 260 (possibly modified by multiplexer 288) unless the branch was caused by a JUMP or JUMP REGISTER instruction, in which case the address specified by the appropriate register will take precedence.

The reason for the specific structure of the branch target address calculating circuits arises from the way the branch target addresses are calculated from the different types of branch instructions, namely a regular branch, JUMP, and JUMP REGISTER. For a regular branch instruction, the relative delay register value, the ATAG register value, and the offset value are added together to create the branch target address; for a JUMP instruction, the ATAG and REL DLY register values are added, and the offset value is concatenated to the sum; and for a JUMP REGISTER instruction, the register value from communication path 296 is used for the branch target address.

The values from ATAG register 227 and RELDLY register 254 are also communicated to a return address adder 300. Return address adder 300 is used to calculate the return address when a branch results in the execution of a subroutine. After the subroutine is finished, it is desirable to return to the instruction immediately following the instruction which called it. Thus, return address adder 300 adds +1 to the addition of the tag, index, and relative delay to produce the address of the instruction following the delay slot of the branch instruction which called the subroutine. The return address is output on a communication path 304.

Fig. 9 shows the Execute (E) and Writeback (W) stages of the pipeline. The contents of ATAG register 227 are communicated to an ETAG register 318 over a communication path 308 and to a compare circuit 341 over a communication path 309, the contents of RELDLY register 254 are communicated to an E REL DLY register 322 over a commu-

5 nication path 312, the calculated return address from return address adder 300 is communicated to a link value (LNVAL) register 326 over communication path 304, and the selected branch target address from multiplexer 289 is communicated to a BR TARG register 330 over communication path 314. An EPC register 334 stores the real address of the instruction the program is supposed to execute in the E stage, and an ASID register stores the program-controlled ASID of the instruction to be executed together with a coherence value (c[2:0]) which typically indicates whether the data used by the instruction is cacheable or not.

10 The ASID and tag stored in ETAG register 318 (corresponding to the instruction fetched) are compared to the ASID and tag from ASID register 338 and EPC register 334 (corresponding to the instruction that is actually supposed to be executed) by a compare circuit 339 to determine if the actual instruction expected to be executed (where the program should be) is actually the instruction fetched from the instruction cache. If the values do not match, then an instruction cache miss signal is provided on a communication path 340.

15 At this time, the value in ATAG register 227 corresponds to the line containing the predicted branch target instruction, whereas the value in BR TARG register 330 corresponds to the actual branch target address. Thus, the index and destination field (the predicted branch target address) received from ATAG register 227 over communication path 309 is compared to the calculated branch target address received from BT TARG register 330 over a communication path 343 by a compare circuit 341 to determine whether the actual branch target instruction expected to be executed corresponds to the predicted branch target instruction fetched from the instruction cache. If the values do not match, then a branch cache miss (branch misprediction) signal is provided on a communication path 345.

20 The value in EPC register 334 is communicated to a WPC register 354 in the writeback stage of the pipeline and to one input of a multiplexer 362 over a communication path 358. The other input to multiplexer 362 receives the value in WPC register 354 (the original value of EPC register 334 delayed by one cycle) over a communication path 366. Multiplexer 350 selects one of these values and communicates the selected value to one input of an EPC adder 350. EPC adder 350 is responsible for updating the value from EPC register 334 during normal operation. The value of EPC register 334 ordinarily is selected during normal operation, and the value of WPC register 354 is selected for exception processing.

25 The other input to EPC adder 350 is coupled to a multiplexer 366. One input to multiplexer 366 is the number of valid instructions dispatched from dispatch register 18, and the other input is an exception adjustment value 369 (-1 to +3). During normal operation, the value from EPC register 334 is incremented by the number of valid instructions dispatched from dispatch register 18 so that the value in EPC register 334 points to the next instruction to be executed. When an exception occurs (trap, instruction cache miss, etc.), the exception adjustment value is added to the value in WPC register 354 to indicate the instruction which caused the exception. The value -1 is used when the exception was caused by a delay instruction, since in that case it is desirable to point to the branch instruction immediately before it. The value indicating which instruction caused the exception is stored in an EPC-COP register 370, which is reloaded with its present value until another exception occurs via multiplexer 374. A TRAP-BASE register 376 stores an address that the program should go to when an exception occurs and communicates the value to a multiplexer 377. The other input to multiplexer 377 is a reset vector value. One of these values is selected and output on a communication path 379.

30 A multiplexer 380 receives the value from EPC-COP register 370 over a communication path 384 when returning from an exception, a vector address from communication path 379 on an exception condition, the calculated branch target address over a communication path 388 for branches, the EPC value from communication path 358 to hold the EPC value during an instruction cache miss, and the updated EPC value over communication path 396. The selected value is output on a communication path 430 (PC JAM BUS[47:0]), of which bits [13:2] are the correction values supplied to the F stage circuitry shown in Fig. 7 to correctly index the instruction cache, tag memory 34A and branch prediction cache 114.

45 During normal operation, the updated EPC value is selected by multiplexer 380 and loaded into EPC register 334. When a branch cache miss occurs, multiplexer 380 selects the calculated branch target address and communicates the new branch target address to branch cache 114 via communication path 160 (Fig. 7). The write address used to update branch prediction cache 114 is calculated by a branch cache address adder 400 which adds the value in EPC register 334 to the value in E REL DLY register 322 and produces the write address on a communication path 404. It should be noted that the value of bits [3:2] on communication path 404 correspond to the position of the delay instruction and may be used as the source field in the branch prediction cache entry. The remaining write data on communication path 168 comprises bits [13:2] of the calculated branch target address, which is the updated index and destination field entries.

Industrial Applicability

55 While the above is a description of a preferred embodiment of the present invention, various modifications may be employed yet remain within the scope of the present invention. Consequently, the scope of the invention should be ascertained from the appended claims.

how many instructions are previously dispatched, and without regard to code alignment or instruction type, and therefore instruction dispatch operation can be performed smoothly so that the execution speed is improved.

Also, branches are predicted at the time of instruction fetch, and the predicted target instruction is fetched immediately so that the target instruction is available for execution immediately after the branch instruction is executed, and therefore the decrease in the execution speed due to the branch instruction minimized.

Claims

1. An apparatus for processing instruction in a computing system comprising:

first and second instruction storing circuits, each instruction storing circuit storing N instructions for parallel output;

an instruction dispatch circuit, coupled to the first instruction storing circuit, for dispatching L instructions stored in the first instruction storing circuit, wherein L is less than N; and

an instruction loading circuit, coupled to the instruction dispatch circuit and to the first and second instruction storing circuits, for loading L instructions from the second instruction storing circuit into the first instruction storing circuit after the L instructions are dispatched from the first instruction storing circuit and before further instructions are dispatched from the first instruction storing circuit.

2. The apparatus according to claim 1 wherein the instruction loading circuit loads the L instructions from the second instruction storing circuit into the positions previously occupied by the L instructions dispatched from the first instruction storing circuit.

3. An apparatus for processing instructions in a computing system comprising:

an instruction storing circuit for storing N instructions for parallel output;

an instruction dispatch circuit, coupled to the instruction storing circuit, for dispatching L instructions stored in the first instruction storing circuit, wherein L is less than N; and

an instruction queue for storing M lines of N instructions from an instruction memory;

an instruction loading circuit, coupled to the instruction storing circuit and to the instruction queue, for loading L instructions from the instruction queue into the instruction storing circuit after the L instructions are dispatched from the instruction storing circuit and before further instructions are dispatched from the instruction storing circuit.

4. The apparatus according to claim 3 wherein the instruction loading circuit loads the L instructions from the second instruction storing circuit into the positions previously occupied by the L instructions dispatched from the first instruction storing circuit.

5. The apparatus according to claim 3 wherein the instruction dispatch circuit comprises a dispatch pointer for storing a value indicating a location of a next instruction to be dispatched in the first instruction storing circuit.

6. The apparatus according to claim 5 wherein the dispatch pointer comprises a modulo-N counter.

7. The apparatus according to claim 6 wherein the instruction queue comprises a queue pointer for storing a value indicating a location of a next instruction to be loaded from the instruction queue into the instruction storing circuit.

8. The apparatus according to claim 7 wherein the instruction queue further comprises a multiplexer, coupled to the queue pointer, for selecting N instructions from the queue and outputting the N selected instructions to the instruction storing circuit.

9. The apparatus according to claim 8 wherein the multiplexer selects the N next sequential instructions from the queue pointer value.

10. The apparatus according to claim 9 further comprising a queue loading circuit for simultaneously loading N instructions from the instruction memory into the instruction queue.

11. The apparatus according to claim 10 further comprising:

a clock for providing periodic clock pulses; and

wherein the queue loading circuit simultaneously loads N instructions from the instruction memory into an empty line in the instruction queue upon even clock pulse.

12. The apparatus according to claim 10 wherein the queue loading circuit includes a load inhibiting circuit for inhibiting loading of instructions from the instruction memory into the instruction queue when there are no empty lines in the queue.
- 5 13. The apparatus according to claim 3 wherein the instruction queue includes an input communication path, and further comprising a queue bypass circuit coupled to the input communication path and to the instruction storing circuit for directly communicating instructions from the input communication path to the instruction storing circuit.
- 10 14. An apparatus for processing instruction branches in a computing system comprising:
an instruction memory for storing a plurality of lines of a plurality of instructions; and
a branch memory for storing a plurality of branch prediction entries, each branch prediction entry containing information for predicting whether a branch designated by a branch instruction stored in the instruction memory will be taken when the branch instruction is executed.
- 15 15. The apparatus according to claim 14 wherein each branch prediction entry includes a branch target field for indicating a target address of a line containing a target instruction to be executed if the branch is taken.
16. The apparatus according to claim 14 wherein each branch prediction entry includes a single-bit branch prediction field for predicting whether a branch designated by the branch instruction stored in the instruction memory will be taken when the branch instruction is executed.
- 20 17. The apparatus according to claim 14 wherein each branch prediction entry corresponds to a line in the instruction memory.
- 25 18. The apparatus according to claim 17 wherein each branch prediction entry includes a branch target field for indicating a target address of a line containing a target instruction to be executed if the branch is taken.
19. The apparatus according to claim 18 wherein each branch prediction entry includes a destination field indicating where the target instruction is located within the line indicated by the branch target address.
- 30 20. The apparatus according to claim 18 wherein each branch prediction entry includes a source field indicating where the branch instruction is located within the line corresponding to the target address.
- 35 21. The apparatus according to claim 20 further comprising:
a counter for storing an address value used for addressing the instruction memory;
an incrementing circuit for incrementing the address value in the counter for sequentially addressing the lines in the instruction memory; and
a counter loading circuit for loading the target address into the counter when the branch prediction entry predicts the branch designated by the branch instruction stored in the instruction memory will be taken when the branch instruction is executed.
- 40 22. The apparatus according to claim 21 further comprising an invalidate circuit for invalidating selected instructions in a line addressed by the address value in response to the source field.
- 45 23. The apparatus according to claim 21 further comprising an invalidate circuit for invalidating selected instructions in a line addressed by the address value in response to the destination field.
- 50 24. The apparatus according to claim 21 further comprising an invalidate circuit for (1) invalidating instructions which follow the branch instruction in the line addressed by the address value in response to the source field, and (2) invalidating instructions which precede the target instruction in the line addressed by the target address in response to the destination field.

FIG.1




FIG.2




FIG.3




FIG. 4

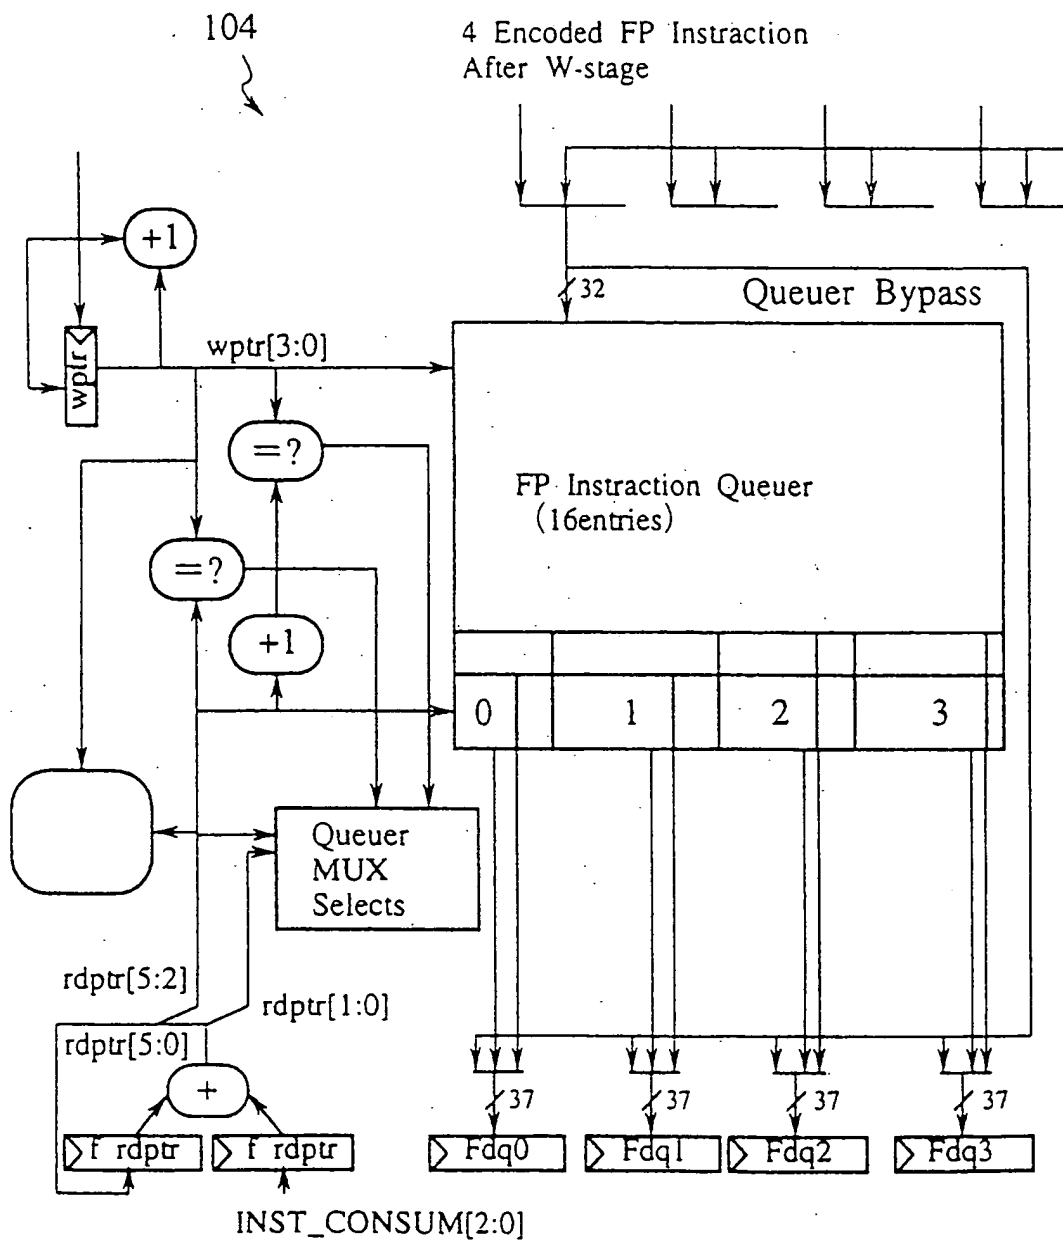




FIG.6





FIG. 7



To Fig. 8

FIG. 8

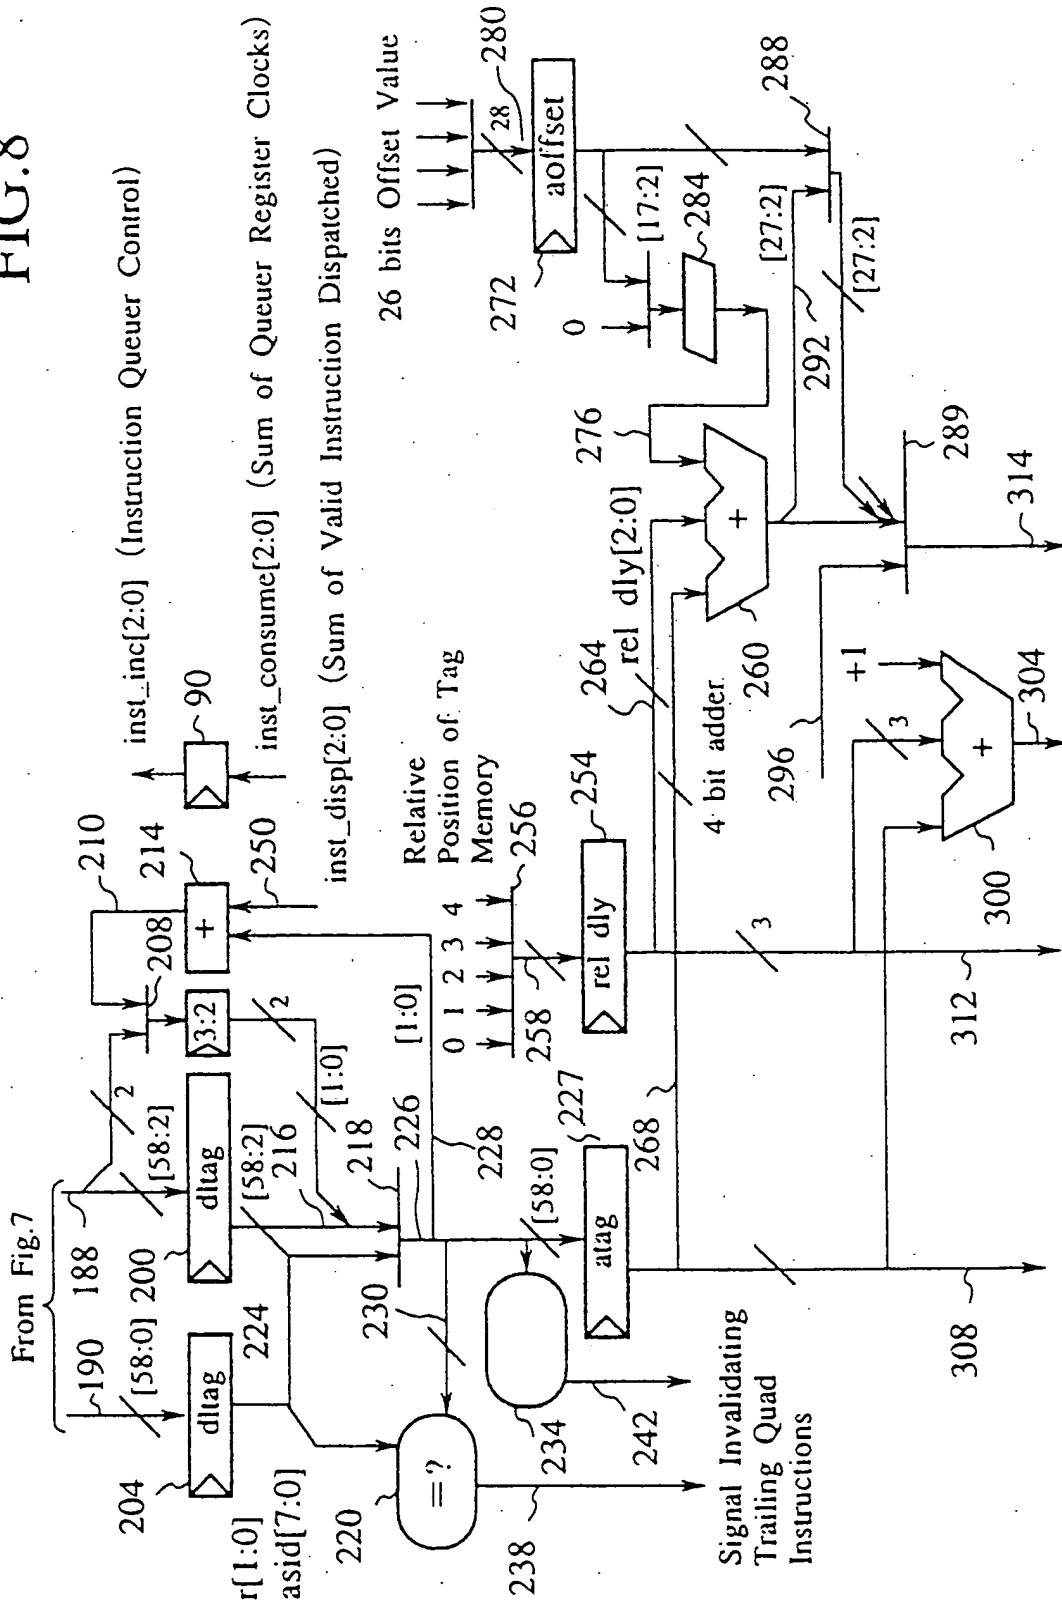


FIG. 9

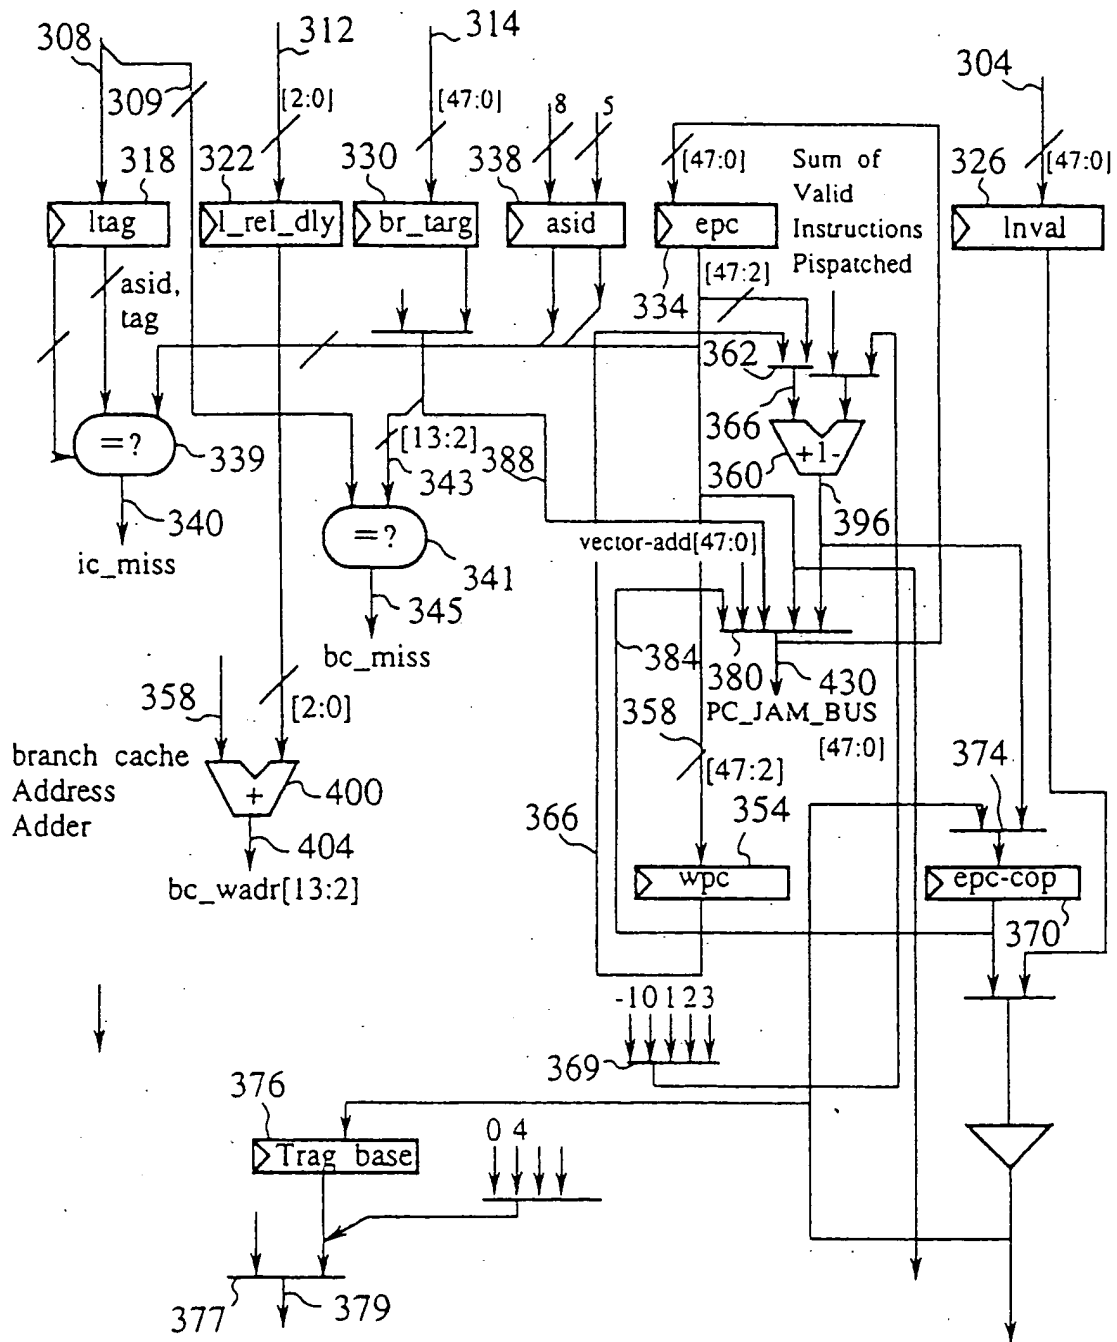



FIG.10



A. CLASSIFICATION OF SUBJECT MATTER		
Int. Cl ⁶ G06F9/38		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols)		
Int. Cl ⁶ G06F9/38		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Jitsuyo Shinan Koho 1926 - 1995		
Kokai Jitsuyo Shinan Koho 1971 - 1995		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	JP, A, 57-29153 (Fujitsu Ltd.), February 17, 1982 (17. 02. 82) (Family: none)	1-13
A	JP, A, 57-152046 (Mitsubishi Electric Corp.) September 20, 1982 (20. 09. 82) (Family: none)	5-12
A	JP, A, 50-152629 (Hitachi, Ltd.), December 8, 1975 (08. 12. 75), Fig. 1 (Family: none)	13
Y	JP, A, 59-91549 (NEC Corp.), May 26, 1984 (26. 05. 84) & US, A, 4984154	14-24
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.		
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier document but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family		
Date of the actual completion of the international search February 15, 1995 (15. 02. 95)		Date of mailing of the international search report March 7, 1995 (07. 03. 95)
Name and mailing address of the ISA/ Japanese Patent Office		Authorized officer
Facsimile No.		Telephone No.